Numerical Weather Prediction

A note on the frequency of radiation calculations in the Unified Model.

Forecasting Research Technical Report No. 485

D.N. Walters, M.R. Willett and S.F. Milton

email:nwp_publications@metoffice.gov.uk

© Crown copyright

A note on the frequency of radiation calculations in the Unified Model.

David Walters, Martin Willett and Sean Milton

8^{th} June 2006

Abstract

We investigate the sensitivity of the Met Office operational global NWP model to more regular long- and short-wave radiation calculations, by comparing a series of forecasts with the radiation schemes implemented both once every hour and once every three hours of model integration time. We observe small differences in objective verification, and a small but systematic increase in the amount of short-wave radiation reaching the surface, which affects the diurnal evolution of the lower atmosphere. We also note that the relative cost of hourly radiation calls in the enhanced-resolution global model (with a 40km grid and 50 vertical levels) is comparable with that of three-hourly calls in the 60km 38 level model.

1 Introduction

In the operational unified model (UM) global NWP model, full calculations of long- and short-wave (LW & SW) forcing are currently performed only once every twelve physics timesteps, i.e. every three hours of simulation time, due to their high computational expense. It is clear, however, that three hours is too long a time between radiation calculations to correctly model all the significant interactions in the atmosphere. For example, in nature, the presence of clouds can lead to a feedback situation: SW warming can lead to thinning of a cloud layer, which will reduce its reflectivity and allow yet more SW radiation to reach the surface. Over a period of three hours, however, the cloud fields in the model can change considerably, such that with such irregular radiative interactions, these secondary effects cannot occur.

With the upgrade of the Met Office's NEC supercomputers and the recent increase in both vertical and horizontal resolutions, the relative expense of the radiation code in an operational UM forecast is decreasing with time; with - or even without - further optimisation of the code, it should be possible to include 24 radiation increments every day, i.e. one every hour of simulation time. Therefore, it is worthwhile investigating what impact this would have on the results of the global forecast.

Similar studies have been performed in the past as part of investigations into NMC problem 182, i.e. the underestimation of winter-time 1.5m temperatures and overestimation of fog and low cloud in the UK and central Iraq. In these tests, it was concluded that the source of the excessively cold boundary layer did not lie in the radiation timestepping, i.e. repeating the

forecasts with 72 radiation timesteps per day was not seen to improve low cloud predictions. Eventually, these problems were attributed to the boundary layer scheme and were seen to improve with the implementation of the HadGEM physics upgrade in January 2005 [1]. It is possible, however, that using only 8 radiation timesteps per day could be detrimental to forecasts in general, which is why we opt to investigate this issue in detail.

In section 2, we outline how the UM deals with model timesteps shorter than radiation timesteps and the resulting consequences. In section 3, we show that the relative computational expense of 24 calls of the radiation routines in the planned higher-resolution model is comparable to that of the three-hourly calls in the current operational model. In section 4, we highlight some large-scale differences noted over the course of a series of 5 day forecasts. Finally, in section 5, we show the impact that the differences discussed in section 2 can have on the diurnal cycle. In particular, we illustrate one example in which the feedback situation described above is seen to occur.

2 Timestepping the UM radiation scheme

It is not the purpose of this document to describe the workings of the radiation scheme in detail; this is outlined in UM documentation paper 23 [2]. In this section, however, we discuss the way in which the scheme deals with having different radiation and physics timesteps, and known problems associated with this. In general, the length of the radiation timestep Δt_{rad} must be a multiple of the physics timestep Δt_{phys}, with this multiple usually chosen to be greater than one. Whilst this means that there are many physics timesteps for which the radiation scheme does not perform a full calculation, the scheme still provides increments in atmospheric temperature T every Δt_{phys}.

In the LW, persistence is used such that the increments calculated at the beginning of the radiation timestep are provided identically over all of Δt_{rad}. As the primary source of LW in the atmosphere is blackbody radiation, this is equivalent to assuming - at least for the purpose of the radiation scheme - that T at the surface, or in a slab of atmosphere, remains constant for the entire radiation timestep. Figure 1 shows a one day time-series of T increments output by the LW radiation scheme for all levels over a single grid-point. The top two panels show results from simulations with the radiation scheme called every three hours (CONTROL) and one hour (referred to as "HOURLY_RAD" from hereon) respectively; the bottom panel shows differences between the two time-series[1]. The results show, as one would expect, that the time-series from CONTROL looks fairly "blocky" on the scale of one day. HOURLY_RAD shows that reducing Δt_{rad} to 1 hour considerably reduces this blockiness. Slightly more important, however, is the issue of temporal sampling. At both 12:00 and 15:00, HOURLY_RAD shows the signal of a convective event. Although convection diagnostics show that such events occur in CONTROL at the same time, these are completely missed in the LW T increments as they do not occur at the beginning of a radiation timestep. Reducing the ratio between Δt_{rad} and Δt_{phys} will not resolve this problem, but should at least reduce its regularity.

The primary source of atmospheric SW radiation is light from the sun, which will clearly change considerably over the period of a few hours, typical of Δt_{rad} in a forecast run. In the UM radiation scheme, it is assumed that all radiation transfer occurs vertically upward

[1] For details of the simulations performed see sec. 5.

or downward within individual atmospheric columns. With $\Delta t_{rad} = \Delta t_{phys}$, the amount of incoming SW at the top of a column is given by

$$SW_{in}(x,t) = S(t) \langle \cos \zeta(x) \rangle_{\Delta t} , \tag{1}$$

where $S(t)$ is the incoming solar radiation from the direction of the sun and $\langle \cos \zeta(x) \rangle_{\Delta t}$ is the cosine of the solar zenith angle ζ - i.e. the angle between the sun and the local vertical - averaged over the timestep[2]. When $\Delta t_{rad} \neq \Delta t_{phys}$, the flux at the top of a column is approximated by

$$SW_{in}(x,t) = S(t) \langle \cos \zeta(x) \rangle_{\Delta t_{phys}} \simeq SW_{in}(x,t_{rad}) \frac{\langle \cos \zeta(x) \rangle_{\Delta t_{phys}}}{\langle \cos \zeta(x) \rangle_{\Delta t_{rad}}}, \tag{2}$$

where t_{rad} is the time in the middle of the current radiation timestep. This is an excellent approximation, since the timescale over which $S(t)$ changes is several orders of magnitude longer than Δt_{rad} and the other dependence on Δt_{rad} is cancelled directly. This is illustrated in Fig. 2, where there is no distinguishable difference between SW_{in} in HOURLY_RAD and CONTROL.

The same approach is used to calculate the temperature increments $\delta T(x,z,t)$. At the beginning of each radiation timestep, the SW radiation scheme calculates $\delta T(x,z,t_{rad})$ using the current atmospheric state and the average ζ over the radiation timestep. What is output to the model each physics timestep, however, is rescaled by the same factor as SW_{in}, i.e.

$$\delta T(x,z,t) \simeq \delta T(x,z,t_{rad}) \frac{\langle \cos \zeta(x) \rangle_{\Delta t_{phys}}}{\langle \cos \zeta(x) \rangle_{\Delta t_{rad}}}. \tag{3}$$

The affect that this has on the model can be seen in Fig. 3, where as with the LW, we show the increments in T for all levels over a single grid-point in CONTROL, HOURLY_RAD and their difference. The most striking result is the "lumpiness" of CONTROL data at the top of the atmosphere (TOA) compared to HOURLY_RAD, and the fact the their difference in this region is clearly periodic, with a period of about 3 hours. To investigate this effect further, we perform simulations using the 38 level aquaplanet model,i.e. an idealised UM configuration with no land surface and a prescribed sea-surface temperature; the model is run with $\Delta t_{phys} = 30$ minutes and $\Delta t_{rad} = 3$ hours.Figure 4 shows the SW T increments on model level 34 for 8 successive physics timesteps, with radiation timesteps starting at 09:30 and 12:30. NB: remember, there is no land in this simulation; the coastlines have been added simply to aid the eye with relative positions. One can see that the distribution is skewed, the skew being to the East at the beginning of the radiation timestep and progressing to the west throughout until the beginning of the next radiation timestep, when the skew jumps back to the East. In the middle of the radiation timestep, i.e. when the value of $\langle \cos \zeta \rangle_{\Delta t_{phys}}$ is closest to $\langle \cos \zeta \rangle_{\Delta t_{rad}}$, the distribution is approximately symmetrical. If we consider the evolution at individual grid-points, rather than the overall distribution, at points near the western-most (leading) edge of the region of SW heating, the temperature increments initially increase with time. However, at the radiation timestep there is a sudden reduction in the

[2]To be precise, by $\langle \cos \zeta \rangle$ we really mean $\langle \max(\cos \zeta, 0) \rangle$, since the incoming solar radiation when $\cos \zeta$ is negative - i.e. when the sun is below the horizon - is zero.

temperature increment, after which it continues to increase as normal. The reverse is true at the eastern-most (trailing) edge of the distribution. This can be seen in Fig. 5, which shows the T increments from both the SW and LW radiation schemes over a grid-point on the equator and on various model levels. In the upper levels, the discontinuities in SW temperature increments on radiation timesteps can clearly be seen. In comparison, Fig. 6 shows that with $\Delta t_{rad} = \Delta t_{phys}$, the upper-level increment profile looks relatively smooth. At lower levels, the increments show a high-frequency noise, which is almost certainly due to the intermittent firing of the convection scheme introducing and removing cloud. The reason for the discontinuities and the resulting skewness of the distribution of increments when $\Delta t_{rad} \neq \Delta t_{phys}$, is that unlike the simple SW flux in Eqn. 2, the dependence of the SW T increments on $\langle \cos\zeta \rangle_{\Delta t_{rad}}$ is non-linear. This means that scaling by the factor in Eqn. 3 does not completely remove the dependence on Δt_{rad} and replace it with a dependence on Δt_{phys}. This is demonstrated in Fig. 7, which shows three idealised evolutions of SW temperature increments with time: $\delta T \propto \langle \cos\zeta \rangle^n$, with $n = \frac{1}{2}$, 1 and $\frac{3}{2}$. The red curve is how this would appear if measured every physics timestep, whilst the blue curve is how it would appear if measured in the middle of the radiation timestep and rescaled on physics timesteps as occurs in the UM. The middle panel shows that when the function is proportional to $\langle \cos\zeta \rangle$ - as with SW_{in} - the two resulting profiles are identical. If the function is sub-linear - as in the top panel - we see a "Batman" shaped evolution, which is very similar to the T increments in the upper levels of Fig. 5. If the function is super-linear - as in the lower panel - we see a "tower-like" evolution, which is similar to the T increments in the lower levels of Fig. 5.

The discontinuities caused by this treatment of the SW scheme are slightly worrying, especially as the SW increments tend to dominate the upper atmosphere energy budget during daytime. Whilst reducing Δt_{rad} to 1 hour will not eliminate this problem, it should reduce it, since the temporal range over which the rescaling takes place is reduced. This in turn will reduce the difference between the ratio of '$\langle \cos\zeta \rangle$'s and unity, such that the approximations to the true evolution of δT should improve.

3 Computational cost

The prime reason for performing radiation calculations only once every three hours and using the methods outlined above at intermediate timesteps is the excessive computational cost of the UM radiation scheme. In this section, we show the results of timing experiments to investigate the cost increase associated with performing one full radiation calculation every hour in the 60km 38 level model. We also show results of a similar experiment using the 40km 50 level global model, which replaced the lower-resolution model in December 2005; the results show that the relative cost of an hourly call of the radiation scheme in the 40km model is similar to that of a three-hourly call in the 60km model.

To measure the performance change at a given resolution, we perform sets of 5 T+120 forecasts, with UM timer diagnostics switched on and calculate the mean CPU and wall-clock times (and associated standard errors) spent on the radiation subroutines (RAD_CTRL, LW RAD and SW RAD) and on the entire forecast. The CONTROL forecasts use UM6.0 (without the HadGEM physics upgrade of January 2005) at the old operational resolution of 432(E-W)×325(N-S)×38 levels, with 8 LW and SW radiation timesteps per period (i.e. per day of model integration time). The test forecasts, HOURLY_RAD, use the same set-up but with 24

LW and SW radiation timesteps per period. The CONTROL and HOURLY_RAD experiments were both performed using 6 PEs of the NEC SX6, with 20 segments used in both the LW and SW radiation schemes and 4 used in the convection scheme. The results from these tests

	CONTROL		HOURLY_RAD	
NEC SX6	CPU time	WC time	CPU time	WC time
Time in radiation:	486(± 29)s	694(± 43)s	1372(± 87)s	1925(±125)s
Total time of run:	3794(±294)s	5149(±659)s	4596(±282)s	6216(±618)s
%age of total:	13(± 1)%	13(± 2)%	30(± 3)%	31(± 4)%

Table 1: Mean CPU and wall-clock times spent in the radiation subroutines in the N216 model.

are listed in Table 1, which shows that the computational expense of the radiation scheme in CONTROL and HOURLY_RAD is about 13% and 30% of the total cost respectively.

Table 2 shows the timings from a similar experiment performed with the (now) operational 640(E-W)×481(N-S)×50 level (UM vn6.1) model. These forecasts have a 1 hour radiation

	$\Delta t_{\rm rad} = 1$ hour (N320L50)	
NEC SX8	CPU time	WC time
Time in radiation:	1137(± 98)s	1884(± 107)s
Total time of run:	9759(±761)s	13359(±1489)s
%age of total:	12(± 1)%	14(± 2)%

Table 2: Mean CPU and wall-clock times spent in the radiation subroutines in the N320 model.

timestep and and 15 minute physics/dynamics timestep and are performed on 8 PEs of the NEC SX8 machine. The results show that in the higher-resolution model, the expense of calling the radiation scheme every hour is comparable with that of CONTROL at about 13% of the total cost.

4 Synoptic timescale effects

4.1 Method

In order to study any wide-scale effects of an hourly radiation call, we perform case studies of five T+120 global forecasts in the summer and five in the winter for both the CONTROL and HOURLY_RAD models outlined in the previous section. The forecasts are started without data assimilation using 12Z start-dumps from the operational archive on the dates listed in Table 3, with each pair of forecasts separated for a period $\gtrsim$14 days in order to minimise the correlation of results due to individual synoptic systems. Once the forecasts are complete, results for various diagnostics are meaned over each season. We also perform objective area

verification against radiosondes, surface observations (SYNOPS) and the operational analysis[3].

Summer	Winter
19/06/2003	17/12/2003
02/07/2003	31/12/2003
15/07/2003	14/01/2004
29/07/2003	28/01/2004
01/09/2003	11/02/2004

Table 3: Dates of the start dumps used in the case study.

4.2 Summary of results

As we have already looked at some time-profiles of SW heating, let us consider this first in looking at meaned data from the case study. Figure 8 shows the northern hemisphere (NH) summer (JJAS) mean of the T+24 zonally meaned SW T increments. The vertical co-ordinate is an estimate of model height, such that the upper levels are emphasised. The main difference can be seen in the uppermost model level (in the upper stratosphere) where there is less heating in HOURLY_RAD than in CONTROL. A similar signal is also visible in the NH winter (DJF). This is a signal that the "Batman" distribution of T increments in the upper atmosphere must overestimate the amount of SW heating there. Figure 9 shows the impact this has on verification against radiosonde observations in JJAS in the NH. At 250hPa, the reduced SW heating at the TOA leads to more of a cold bias in HOURLY_RAD, which is slightly worse than CONTROL . This is offset, however, by a slightly higher temperature at 850hPa, which reduces the negative bias seen in CONTROL. These signals are consistent in the NH, tropics and SH in both seasons. Verification of T against surface observations is better in the NH and tropics in JJAS, but worse in this season in the southern hemisphere (SH) and in all areas in DJF.

The main signal in the wind data - an improvement - is a slight increase in the zonal wind speed at higher latitudes and medium level (400-800hPa), which builds over the forecast period (see the Day 3 results in Fig. 10). There is little difference in the plots or verification of surface winds. In pressure at mean sea-level (PMSL) and 500hPa height, the differences in verification show a slight decrease in bias for most areas. For a typical example see Fig. 11.

One of the more interesting differences between HOURLY_RAD and CONTROL can be seen in Fig. 12, which shows the differences in mean net surface SW flux at T+24 in both seasons. The mean difference appears to be positive, with the highest values off the western coasts of the major continents. In these regions, cool continental air passes over warm ocean and is trapped by descending air from the Hadley circulation to form large decks of strato-cumulus (Sc) cloud. In section 5, we demonstrate that this increase in surface SW flux is due to the thinning of this Sc in HOURLY_RAD. These differences remain throughout the forecast period, although they appear less organised by T+120.

[3]One should bear in mind the fact that when comparing with both HOURLY_RAD and the CONTROL, the analysis cycle is advanced using the control model; for this reason, we concentrate on verification against observations.

5 Diurnal timescale effects

5.1 Method

In order to study the diurnal cycle, we perform a T+36 forecast for both CONTROL and HOURLY_RAD, from the 12Z dump on 19/06/2003, i.e. in the NH summer and the SH winter. Cloud-, radiation- and precipitation-related diagnostics are measured every hour, and the diurnal cycle studied using data from T+9 to T+32.

We analyse the data using two methods: *(i)* we determine the local time for every hour at each grid-point and zonal means are taken of data evaluated at the same local time, rather than the same verification time[4]. This gives us a time-series of zonal means which includes a diurnal cycle, and that is more valid statistically than a single grid-point on a single forecast. It also implies that the data at each local time should contain approximately the same degree of forecast error, as they all contain data from T+9 to T+32. Radiation diagnostics, which are measured only every radiation timestep, are treated slightly differently as the value represents the average over the following radiation timestep. In the control job, therefore, these are copied over the next three hours. In HOURLY_RAD, where we have three measurements available in this period, we take the mean of these three and copy it into the three hour period. We then make difference plots using both the raw data and this 3-hourly meaned data, which allows us to see both the instantaneous differences in radiative heating and the effect integrated over 3 hours to match CONTROL. Whilst this method of analysis gives a fair presentation of data, especially for the radiation-related diagnostics, it does tend to "smear out" some of the signal of a diurnal cycle in the process. Because of this, we also choose *(ii)*, to study time-series over the same period at a selection of model grid-points, listed in Table 4.

Region	Long.		Lat.	
N. Pacific	140°	W	20°	N
S. Pacific(A) (near Polynesia)	120°	W	20°	S
N. America	100°	W	40°	N
S. Pacific(B) (near Peru)	90°	W	10°	S
S. America	60°	W	10°	S
N. Atlantic	40°	E	40°	N
Central Africa	20°	E	10°	N
Central Asia	70°	E	50°	N

Table 4: Grid-points at which diurnal time-series were measured.

5.2 Summary of results

As discussed in the previous sub-section, the process of meaning data valid at the same local time smears out a lot of information. In particular, there is little diurnal signal visible in the

[4]Here, we define "local time" by binning grid-points into 24 bands of longitude 15°wide. This is local time in an astronomical sense, and takes no account of international borders and conventions.

results for temperature, humidity and rainfall. A clear diurnal cycle can be seen, however, in the surface energy budget. Figure 13 shows the net surface SW (down) flux, when zonally averaged at the same local time. The top-right panel shows the difference between HOURLY_RAD and CONTROL when treated as hourly time-series. As the SW flux in CONTROL is constant over the 3 hour timesteps, the predominant signal is visible when the incoming radiation changes most over these periods. When the sun first rises (or is about to set) there is a lower flux in HOURLY_RAD than in CONTROL, whilst the flux is consistently higher in HOURLY_RAD over the 6 hours period around noon. The bottom-right panel shows the difference between HOURLY_RAD and CONTROL when the data in HOURLY_RAD is averaged over the three-hour bins. One can see that the main difference between HOURLY_RAD and CONTROL is an increase in the SW flux, most coherently over the periods covering dawn and dusk. This signal can be seen more cleanly in the clear-sky surface SW diagnostic illustrated in Fig. 14. Here, we see clearly that the effect of more regular radiation calls integrated up to 3 hours - at least in the lower atmosphere - is an increase in SW flux during sunrise and sunset. The fact that there are residual differences between HOURLY_RAD and CONTROL in the net SW flux in the middle of the day, which are not there in the clear-sky diagnostic, must be due to differences in the cloud field. Figure 15 shows the zonally averaged net LW radiation flux at the surface. The diurnal signal in CONTROL is seen to lag behind that of HOURLY_RAD by about $1\frac{1}{2}$ hours, which is as expected since the LW rad increments simply persist for the whole radiation timestep. The shape of the bipolar signal in the differences is closely related to the distribution of land, as the main changes in surface LW fluxes over diurnal timescales are due to the evolution of land-surface temperatures.

When looking at the single grid-point data, the Africa grid-point (for which our period corresponds approximately to a single solar day) highlights one of the differences between the diurnal cycle in CONTROL and HOURLY_RAD "seen" by the model. The surface clear-sky SW down, illustrated in Fig. 16, shows that not only does the diurnal cycle look more realistic in HOURLY_RAD, but that if we average this data over three-hour periods, there is $\approx 6 \mathrm{Wm}^{-2}$ more incoming radiation over the periods covering dawn and dusk (agreeing with the results of the averaged data in Fig. 14). This signal is present at all of the grid-points studied.

Data for the point in the N. Pacific shows what meteorological impact this can have on the model. The Case study data shows that this point is covered by one of the decks of Sc cloud previously discussed, which is reflected in the grid-point large-scale cloud amount illustrated in Fig. 17. Over the 3 hour period following dawn (07:00-10:00), where HOURLY_RAD has a higher SW surface flux, the bulk cloud diagnostic shows that this cloud thins, which should allow even more SW radiation down to the surface. Figure 18 shows that in the following 3 hour period, where the clear-sky SW down in HOURLY_RAD and CONTROL are again the same, the result of this thinner cloud is that the net SW downward flux at the surface is over $150 \mathrm{Wm}^{-2}$ higher in HOURLY_RAD. This explains the wide-spread increase net SW between HOURLY_RAD and CONTROL observed in the case study results (Fig. 12) off the western coast of the major continents. The effect is especially noticeable in the summer hemisphere.

The affect that this can have on other diagnostics can be seen in the short-wave temperature increments (see Fig. 19), where we see slightly less heating in the cloud layer (as there is less absorption) and more in the atmosphere below the cloud. This example serves to show how the important feed-back between the radiation and cloud schemes over the timescale of a couple of hours can lead to large (>40%) differences in SW radiative forcing at the surface, which in

turn will affect other parts of the model, such as the evolution of the boundary layer. This difference in surface energy incident on the ocean below the Sc decks will presumably have a bigger impact when the model is coupled to the ocean as in climate change studies.

5.3 24 hour radiation budget

Finally, we use the data analysed in the previous section to calculate radiation budgets at the top of the atmosphere and the surface, by meaning the fields at these levels both over the globe and over the 24 hour period. The meaned results are presented in table 5.3; the standard deviations of the fields were also calculated, but the two forecasts are so correlated, that this does not give any meaningful measure of the error between them.

Flux	CONTROL	HOURLY_RAD	DIFF
Net SW Down (TOA)	230.67	232.31	+1.64
Clear Sky Net SW Down (TOA)	282.11	282.99	+0.88
Net SW Down (Surface)	152.25	154.30	+2.05
Outgoing LW (TOA)	239.13	239.21	+0.08
Clear Sky Outgoing LW (TOA)	264.41	264.41	0.00
Net LW Down (Surface)	-54.13	-54.09	+0.04
Net LW & SW Down (TOA)	-8.46	-6.90	+1.56

Table 5: Global mean radiation fluxes over 24 hours for CONTROL, HOURLY_RAD and their difference (HOURLY_RAD−CONTROL). All figures are in Wm^{-2}.

The net SW at TOA is higher in HOURLY_RAD than in CONTROL. Since the incoming solar radiation is identical by definition, this means that there is less radiation reflected by the planet, i.e. the planet appears darker from "the outside". The value for clear sky net SW (TOA) shows that about half of the increase in net SW is due to a reduction in cloud, although this diagnostic also has a residual increase in HOURLY_RAD, which is probably due to changes in atmospheric absorption at dawn and dusk. The surface SW agrees with the diagnosis of reduced cloud, as more radiation appears to be reaching the surface. The table also show that over the full 24 hours, the net LW has a much weaker dependence on the radiation timestep, and the clear sky diagnostic shows that what signal there is is largely due to differences in cloud. The significant local differences throughout the diurnal cycle appear to almost cancel each other out when calculating the daily average. Overall, we see that the change from CONTROL to HOURLY_RAD has reduced the residual TOA radiation loss from $8.46Wm^{-2}$ to $6.90Wm^{-2}$, i.e. by almost 20%. We should bear in mind, however, that as this figure is only for 1 forecast of a 24 hour period, this result will have little statistical significance.

6 Conclusions and suggestions for further work

We have outlined systematic problems that can be caused by the temporal sampling of radiation calculations in the UM global forecast. These include "missing" short-lived events (such

as convection) in the persistence approach used in the LW part of the spectrum, and the poor interpolation of SW temperature increments by the use of astronomical insolation information, particularly at sunrise and sunset.

We argue that, whilst not solving these problems altogether, increasing the frequency of sampling from once every three hours to once every hour would improve matters. This is especially true for SW radiation, as currently there are as few as four sunlit radiation calls per day at the equator. Over short time-periods, the increased frequency of radiation calculations leads not only to more realistic looking time-series of radiation-related diagnostics, but to a net increase in surface incoming clear-sky SW at sunrise and sunset. In one case, where we study the evolution in detail, this is shown to thin the cloud in the North Pacific strato-cumulus deck, which in turn allows more SW to reach the surface and the boundary layer. These effects combined are shown to decrease the overall albedo of the globe and bring the TOA LW & SW fluxes closer to equilibrium. In the larger-scale evolution, the results of a 10 forecast case study from summer 2003 and winter 2003/4 show that this is translated into a systematic increase in net surface SW flux off the western coasts of the major continents. We also observe a slight increase in zonal winds, moving the forecast closer to observational results, which could be due to more energy being added to the atmosphere by the SW radiation scheme. The slightly warmer 850hPa temperatures observed are also an improvement, although this is at the cost of a larger cold bias at the top of the atmosphere. In summary, whilst there are only slight differences in the objective verification between HOURLY_RAD and CONTROL, we believe that the improvements show that it is worth increasing the frequency of radiation calls in the operational forecast from three-hourly to hourly. It is also clear that more frequent radiation calls are a necessary, but probably not sufficient condition for the diurnal cycle to be correctly represented in the UM.

We have shown that in an N216 38 level forecast, the extra calculations required to run the full radiation scheme every hour rather than every three hours translate into an increase in the cost of the radiation scheme from 13% to 30% of the total compuatational expense. We have also shown, however, that this figure reduces back to about 13% when we perform hourly radiation calculations in an N320 50 level job, such that the relative cost increase as part of the recent resolution enhancement is approximately neutral. One way to reduce this cost even further might be to split the radiation processes into "fast" absorption (i.e. by ice/water clouds, aerosols or weak gaseous absorbers (minor atmospheric gases)) and "slow" absorption (i.e. absorption by strong gaseous absorbers (major atmospheric gases)), assigning a 1 hour timestep to the former and a 3 hour timestep to the latter.This would lead to the same improved interaction with cloud that we have seen here by simply decreasing the radiation timestep. One would probably not see the same improvements in the clear-sky SW fluxes and heating rates, however, as the main absorption in the upper stratosphere, where the largest differences in SW T increments are observed, are in the Hartley band of ozone, which is classified as a slow absorber. Also, recent studies of the evolution of cloud-free convective mixed layers at the Southern Great Plains "Atmospheric Radiation Measurement" (ARM) site have shown that running with a radiation time-step equal to the physics time-step rather than 3 hours can have a positive influence on boundary layer temperature (BL) evolution. The model's cold bias and lag in evolution in BL height are both improved[3]. This suggests that in clear sky conditions the interaction of radiation with the moisture in the evolving BL can be important. This raises a further question mark over the proposed strategy for saving time in radiation calculations by

only having a shorter radiation time-step for "fast" absorbers. We shall leave the investigation of impact of this new scheme for a later study.

References

[1] Sean Milton, Malcolm Brooks, Adrian Lock, Eoin Whelan, and Damian Wilson. HadGEM1 physics for the global NWP model (cycle G34) Improvements to boundary layer, large scale precipitation, convection and saharan albedo. Technical Report 458, NWP, Met Office, Exeter, 2005.

[2] John Malcolm Edwards, Jean-Claude Thelen, and William James Ingram. *The radiation code*. UM documentation paper N^o 23, Met Office, Exeter, 2004.

[3] Sean Milton. Personal Communication.

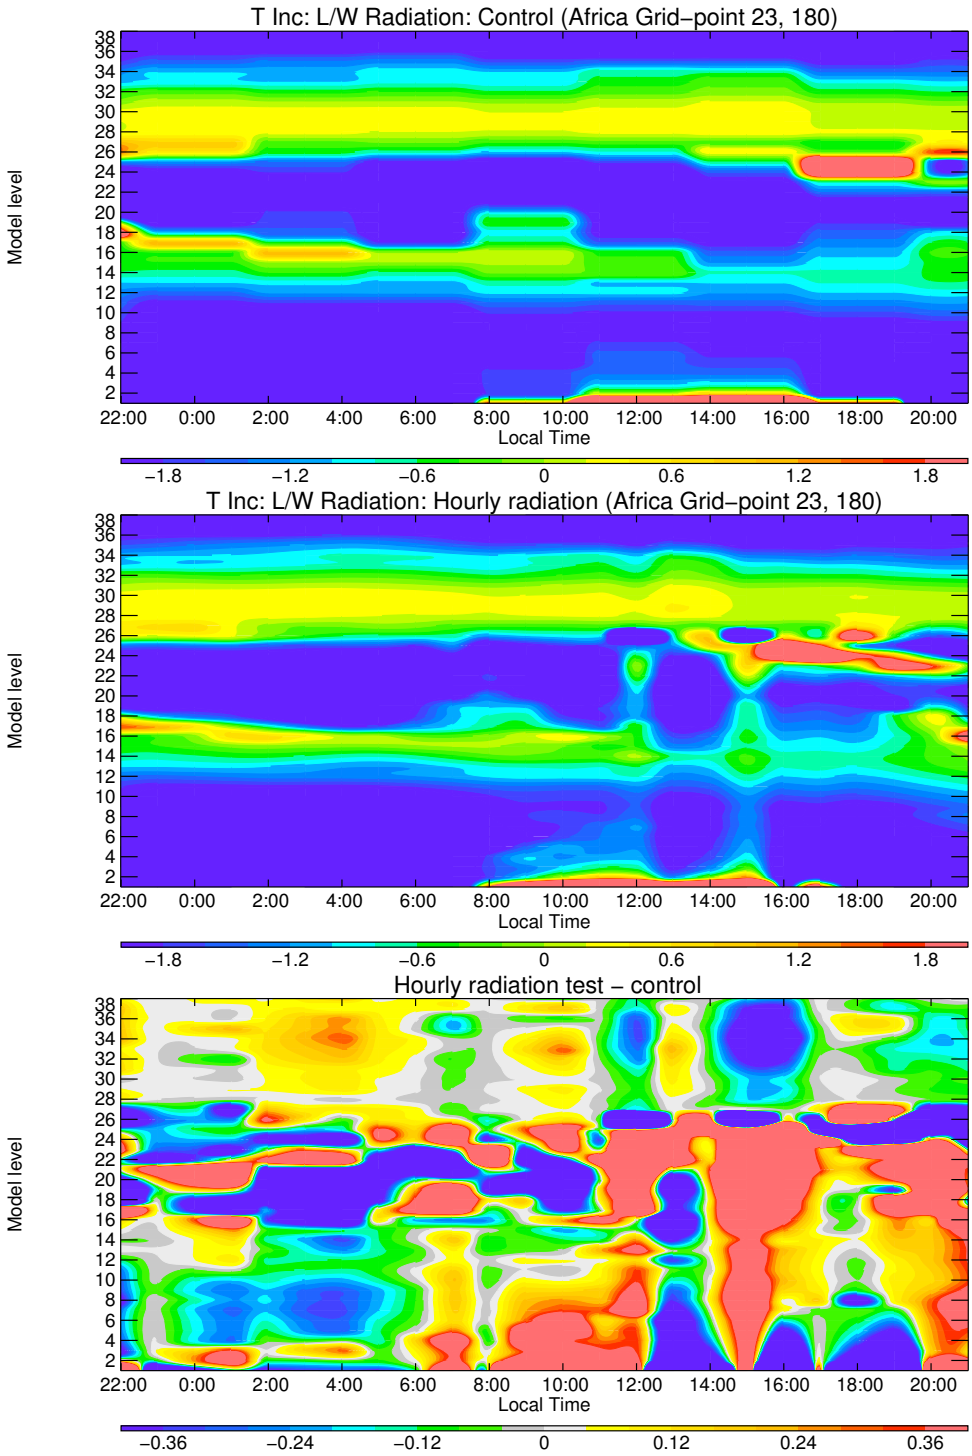

Figure 1: Time-series of T increments (K/day) from the LW radiation scheme for a grid-point in Africa.

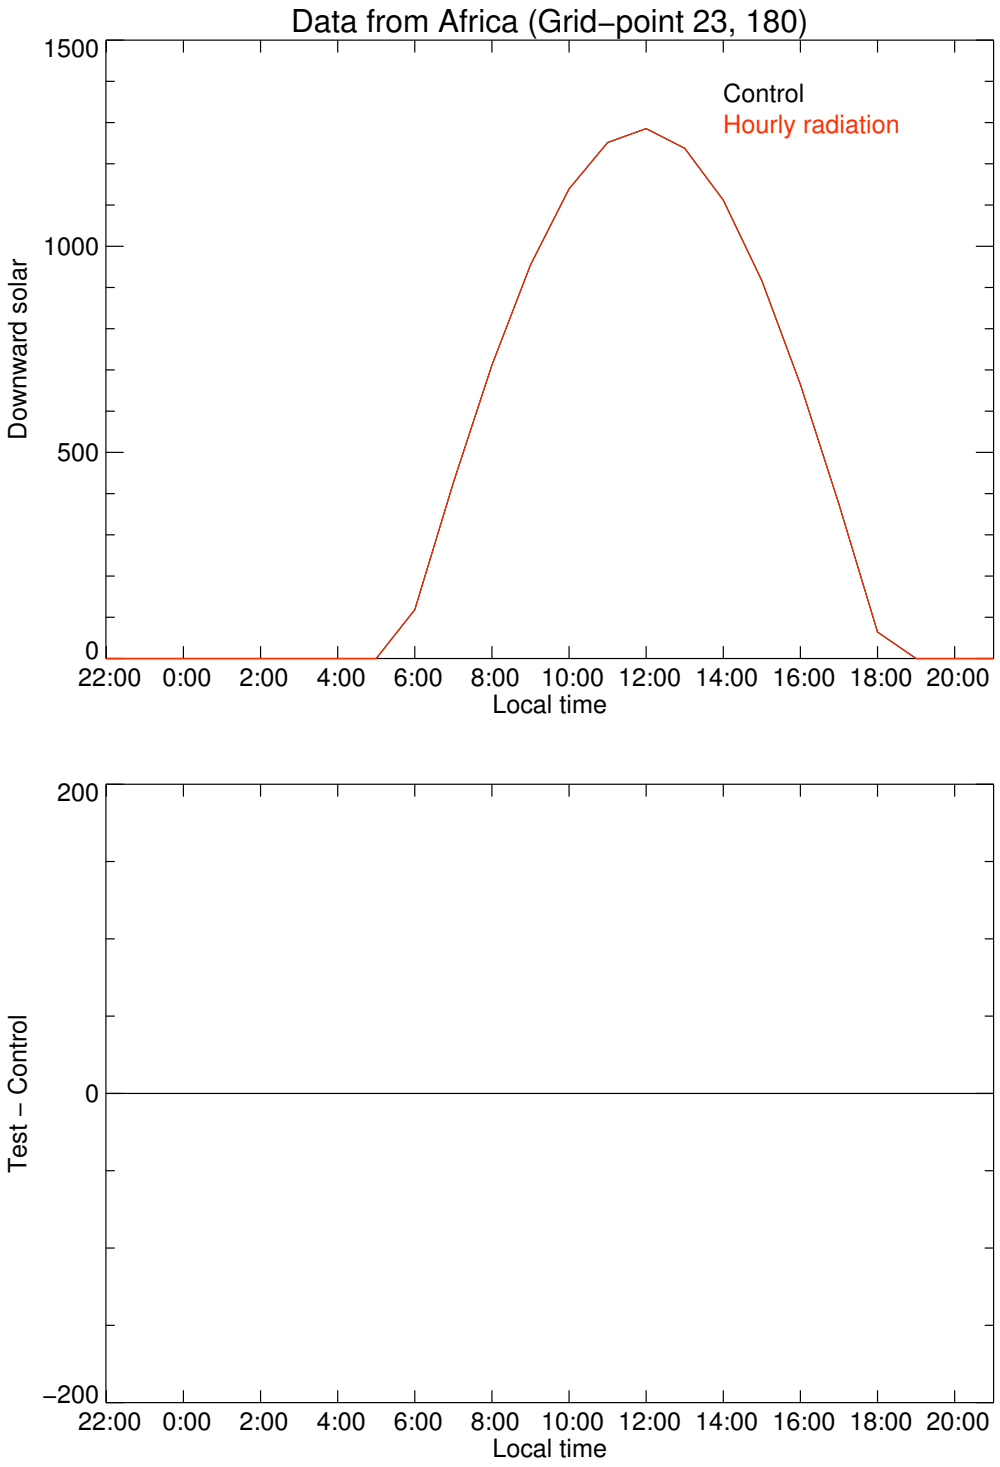

Figure 2: Time-series of incoming SW at TOA for a grid-point in Africa.

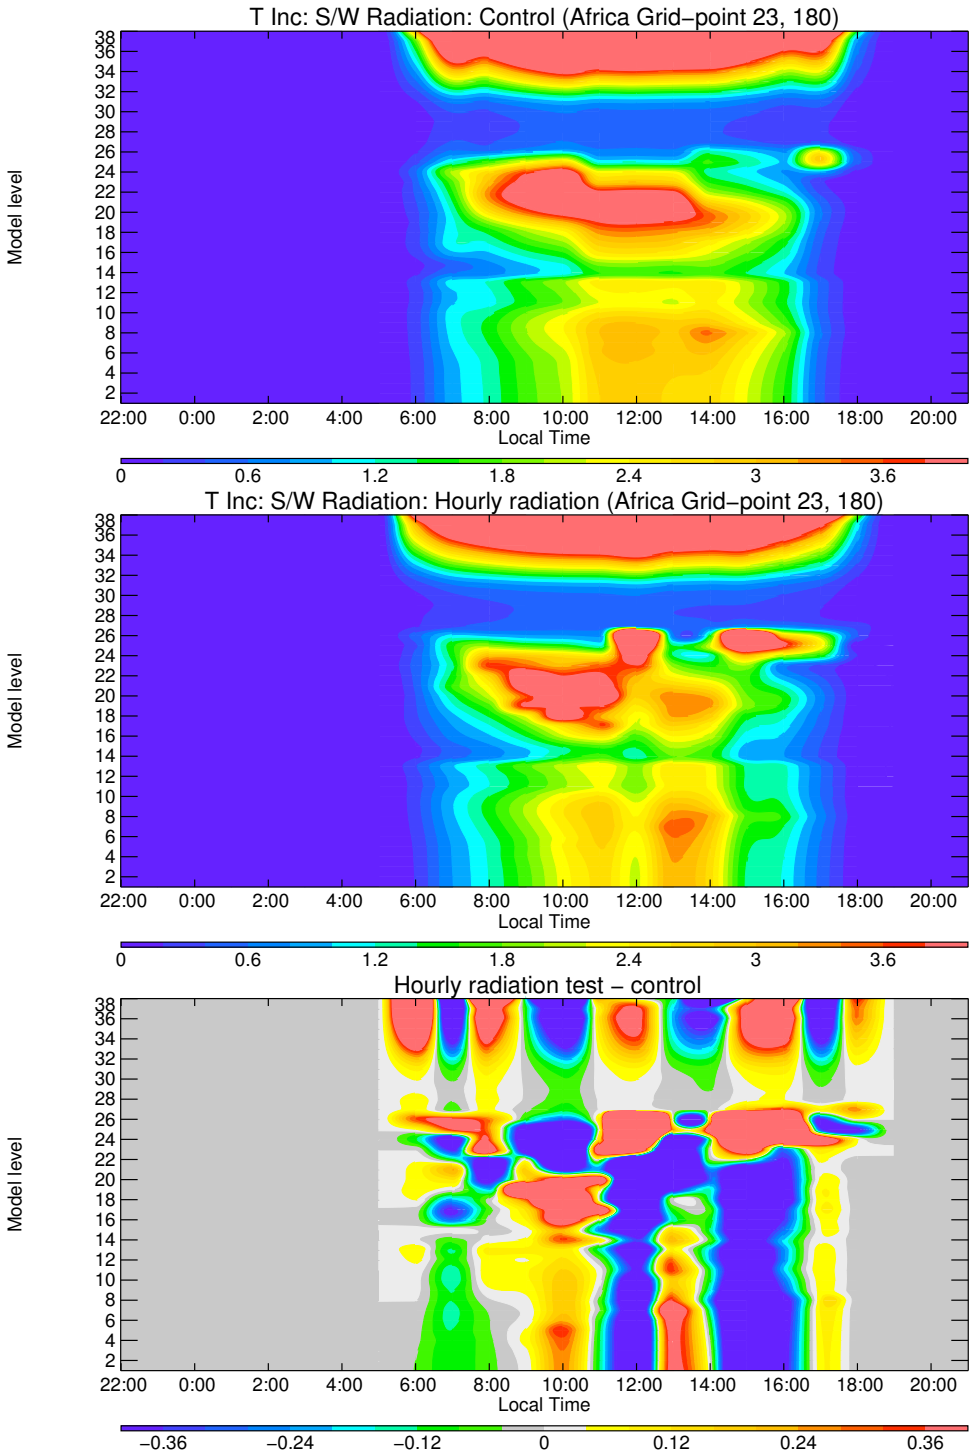

Figure 3: Time-series of T increments (K/day) from the SW radiation scheme for a grid-point in Africa.

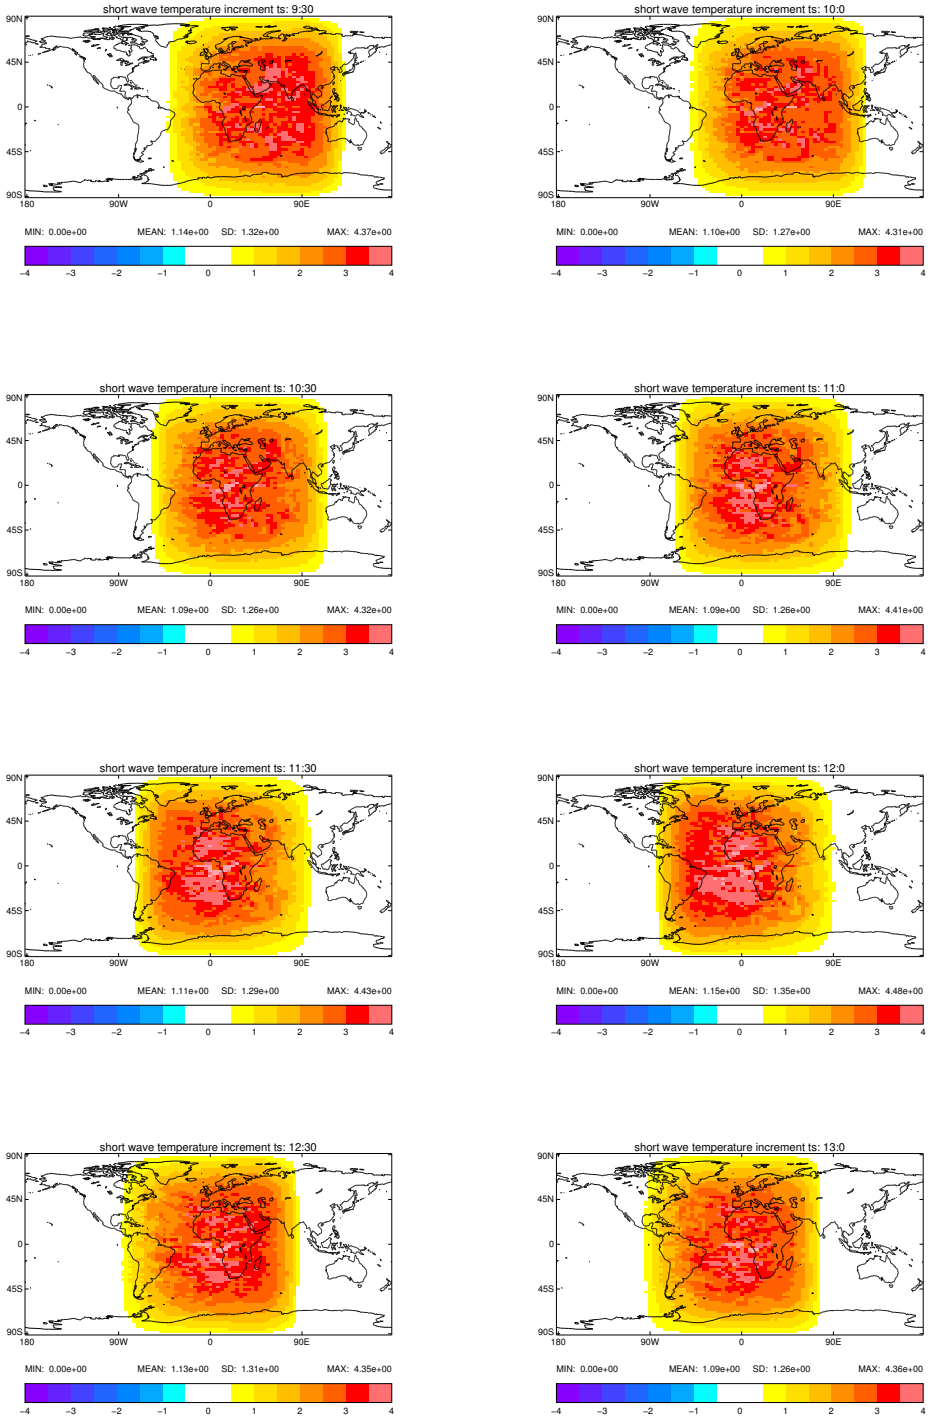

Figure 4: SW T increments (K/day) on model level 34 at successive physics timesteps in the aquaplanet model.

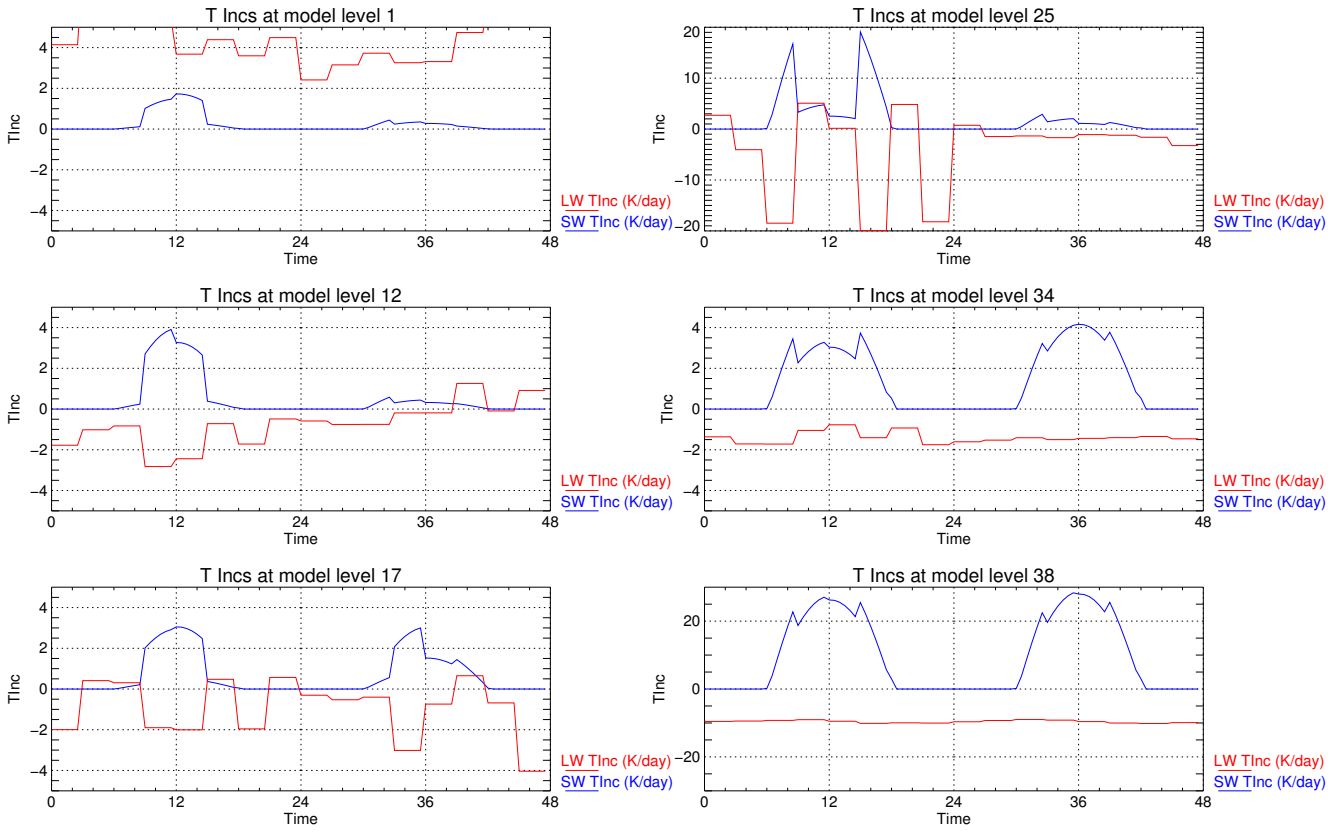

Figure 5: Time-series of T increments (K/day) for several model levels over a single grid-point of the aquaplanet model with $\Delta t_{rad} = 6\Delta t_{phys}$.

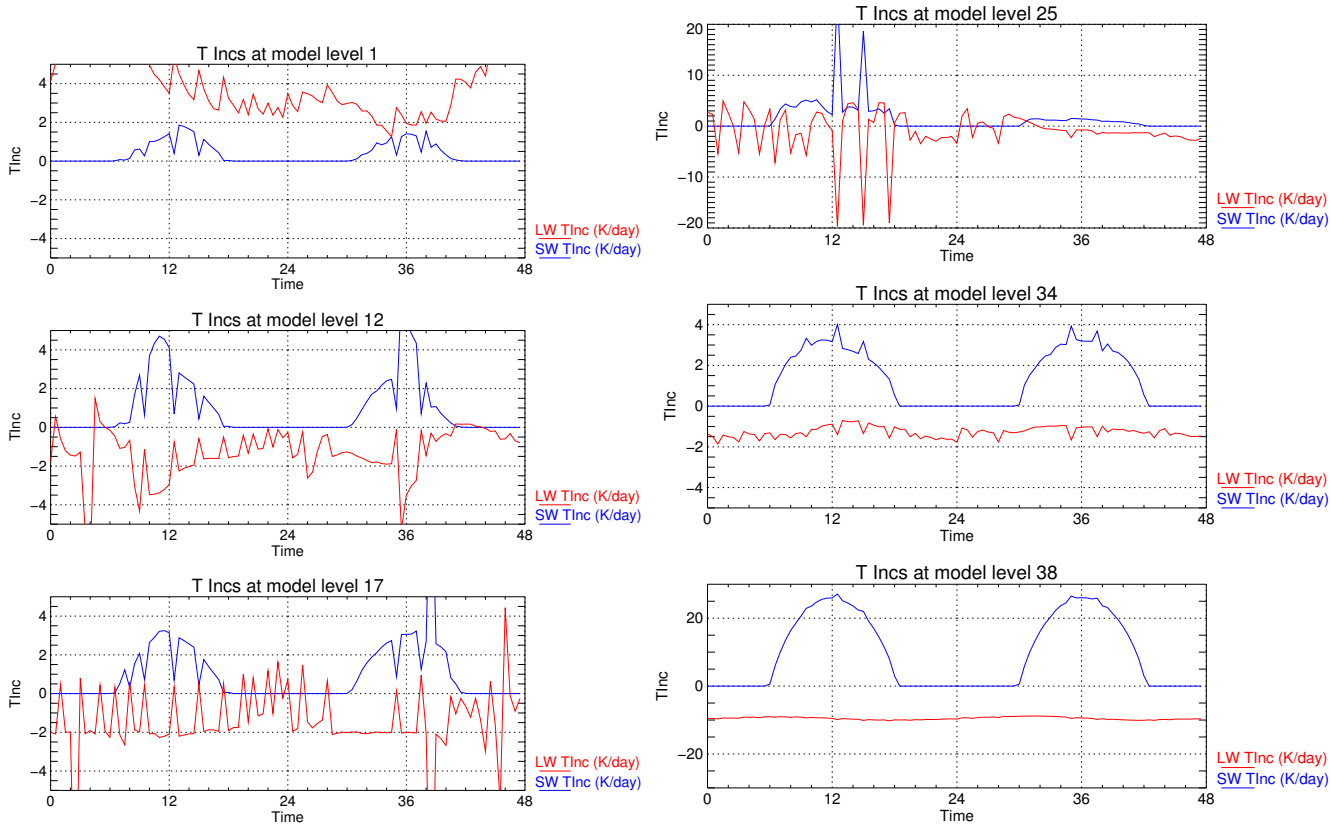

Figure 6: Time-series of T increments (K/day) for several model levels over a single grid-point of the aquaplanet model with $\Delta t_{rad} = \Delta t_{phys}$. The source of the high-frequency noise is intermittently parametrised convection.

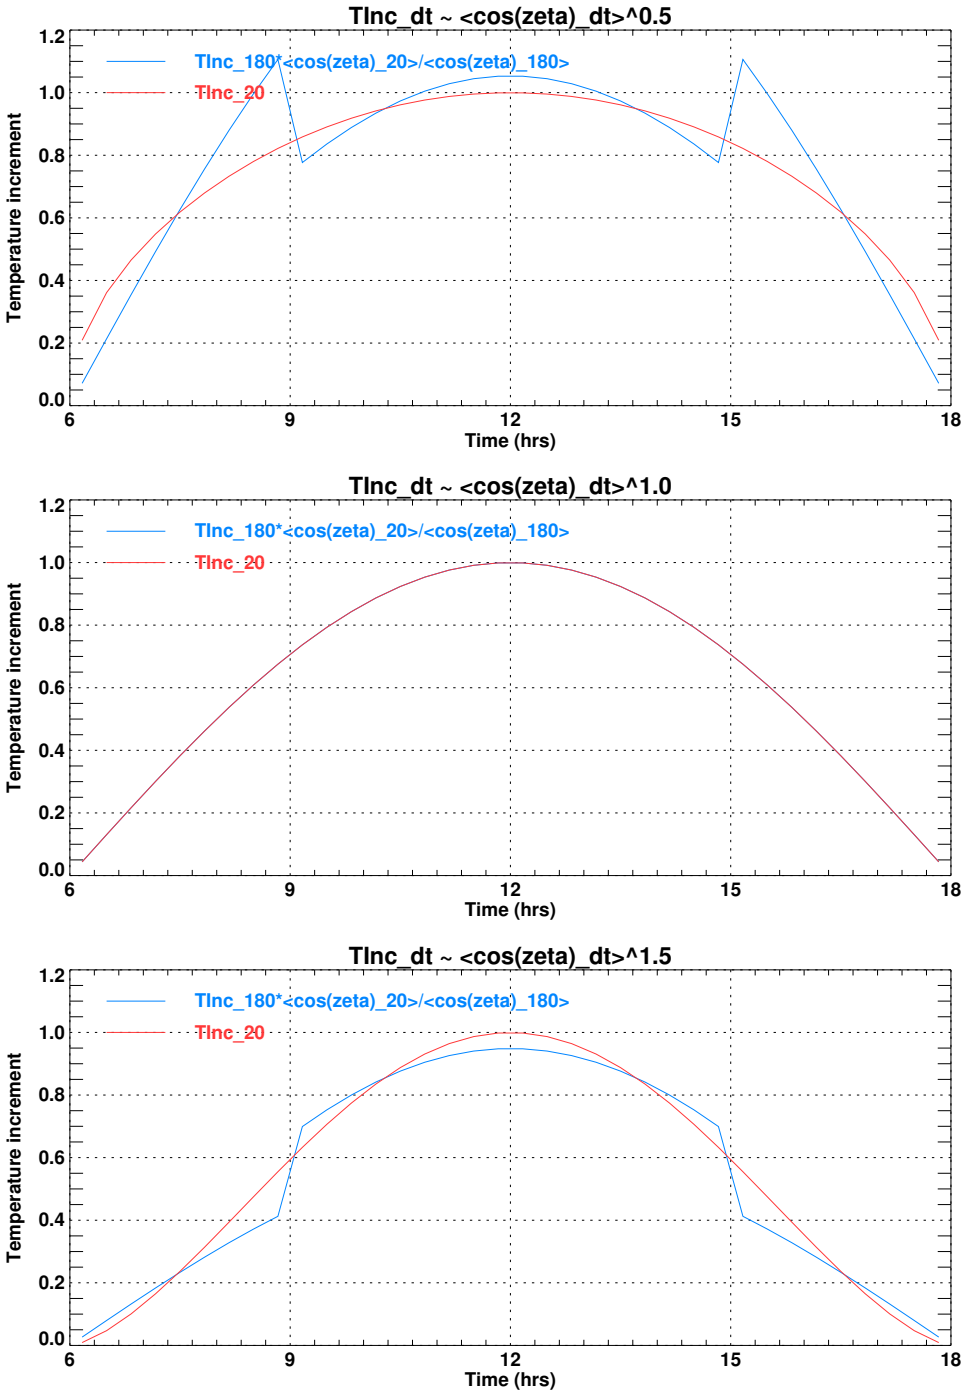

Figure 7: Three idealised developments of SW T increments with $\Delta t_{rad} = \Delta t_{phys} = 20$ minutes (red), and how this would appear to the model with $\Delta t_{rad} = 3$ hours (blue).

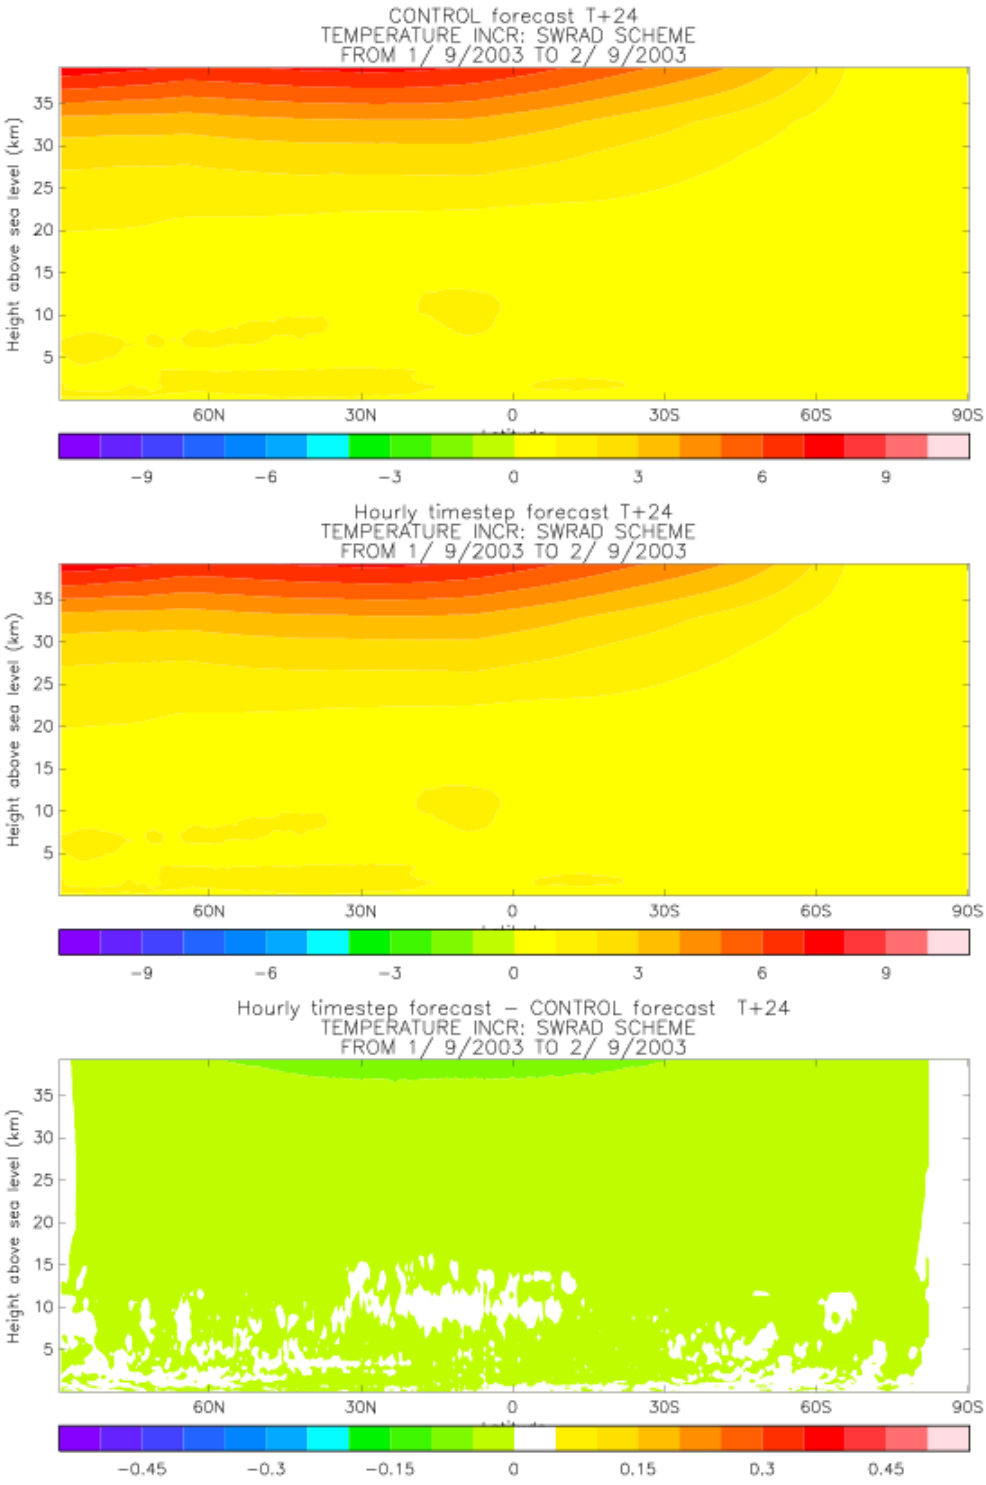

Figure 8: Zonal mean of T+24 SW T increments (K/day) from JJAS season of the case study.

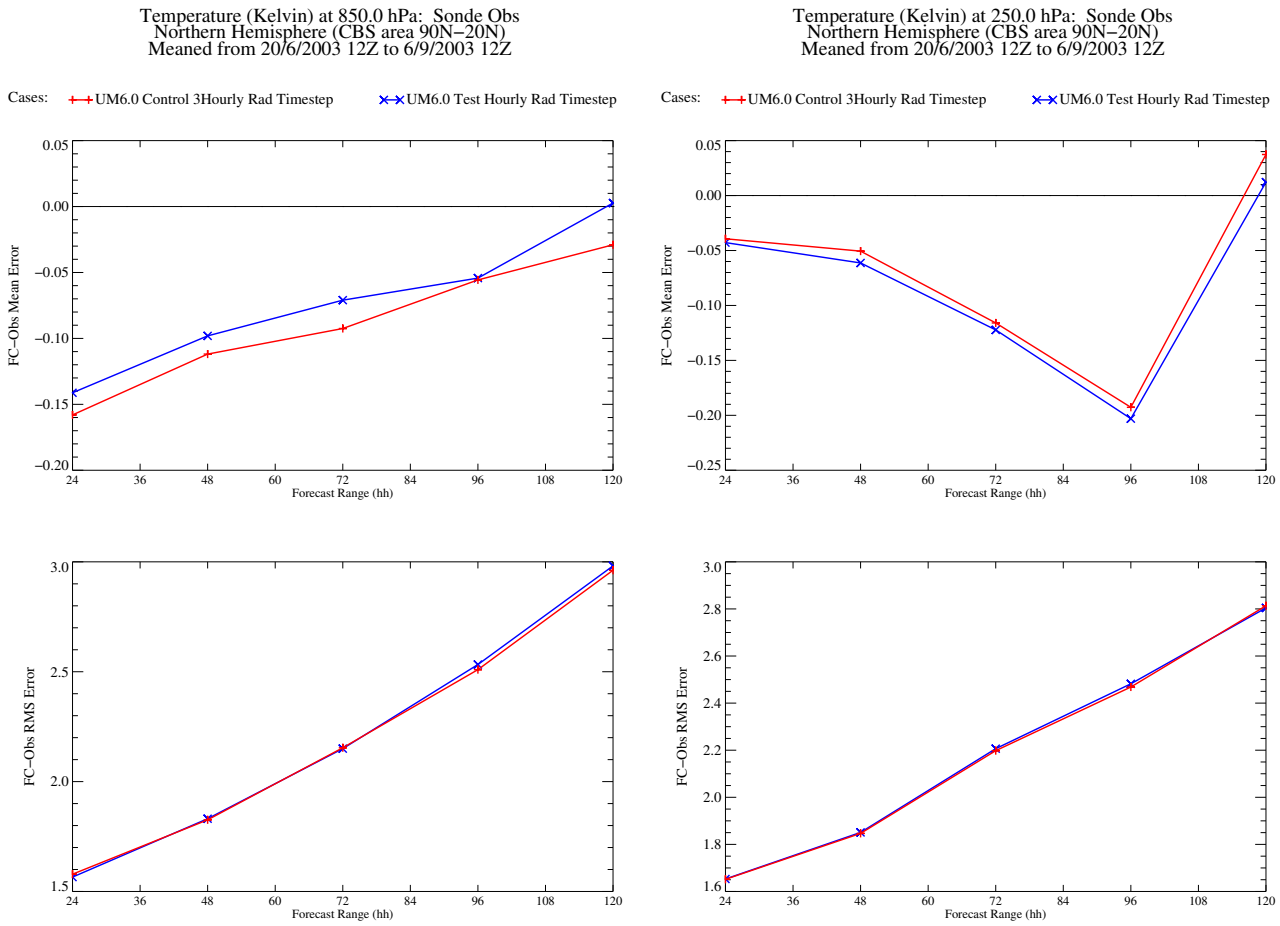

Figure 9: NH JJAS verification of T vs. radiosonde observations at 850 and 250hPa.

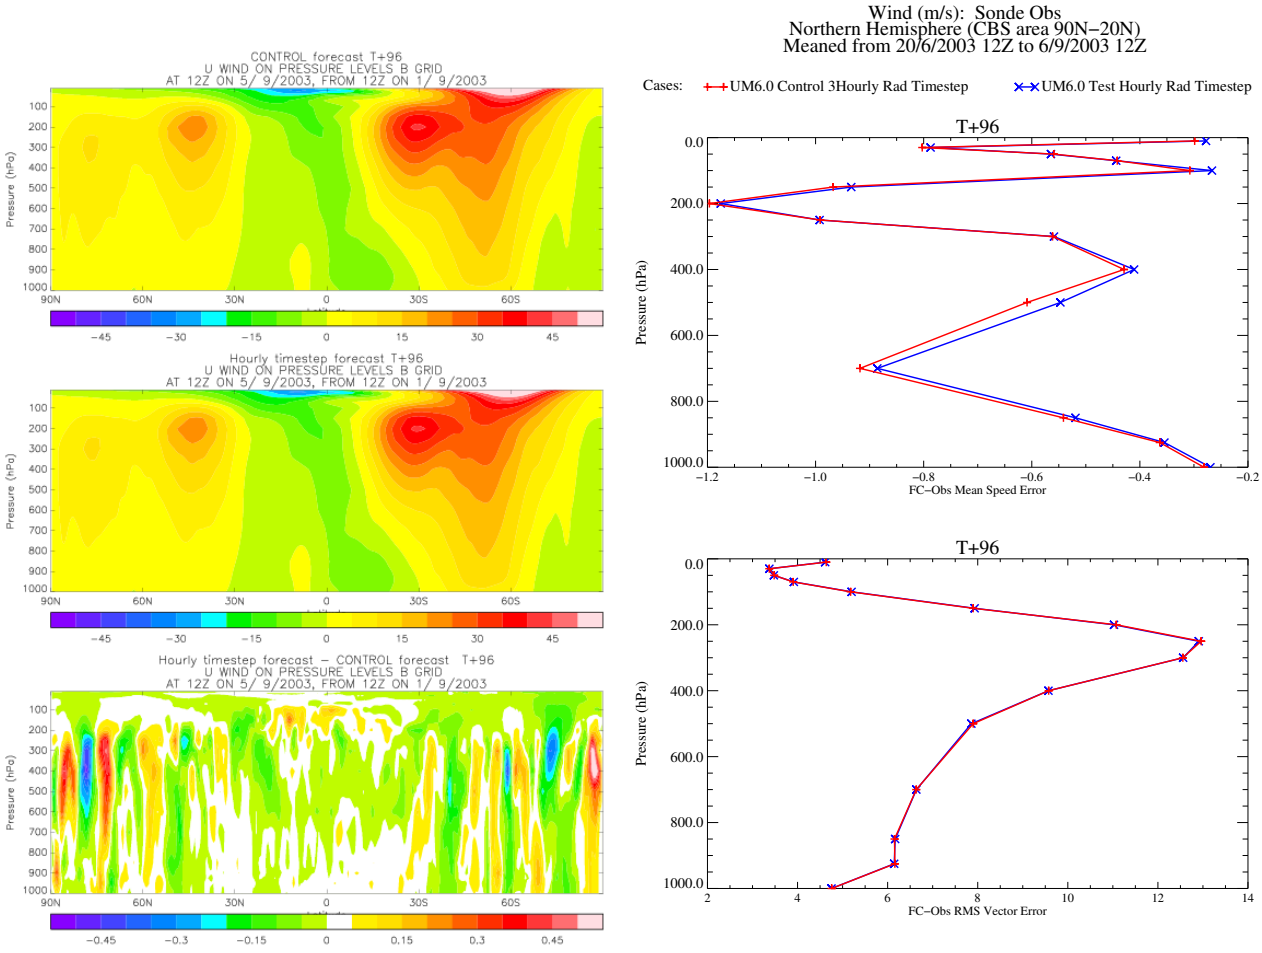

Figure 10: Day 4 zonal wind and NH verification against radiosondes in JJAS.

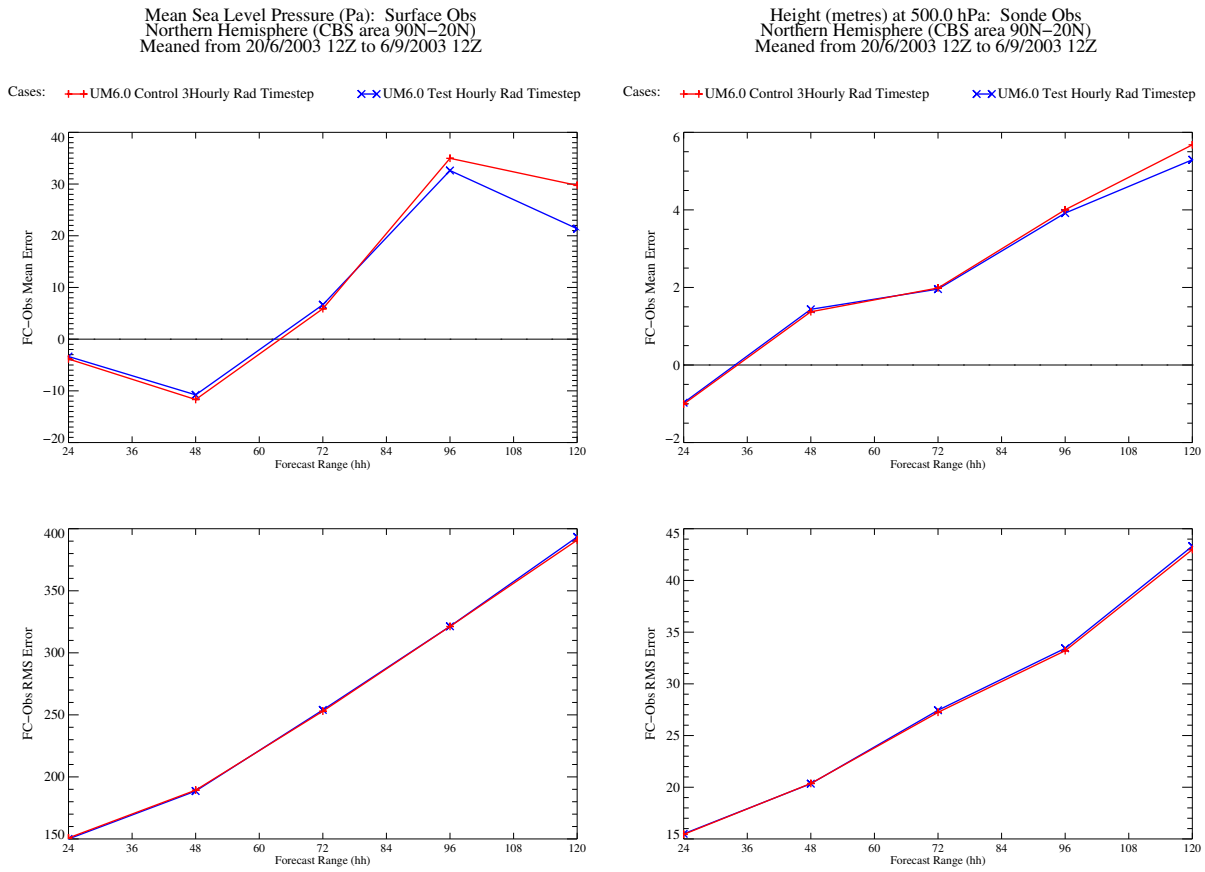

Figure 11: Mean sea level pressure and 500hPa height verifications against SYNOPS and radiosondes respectively in the NH summer.

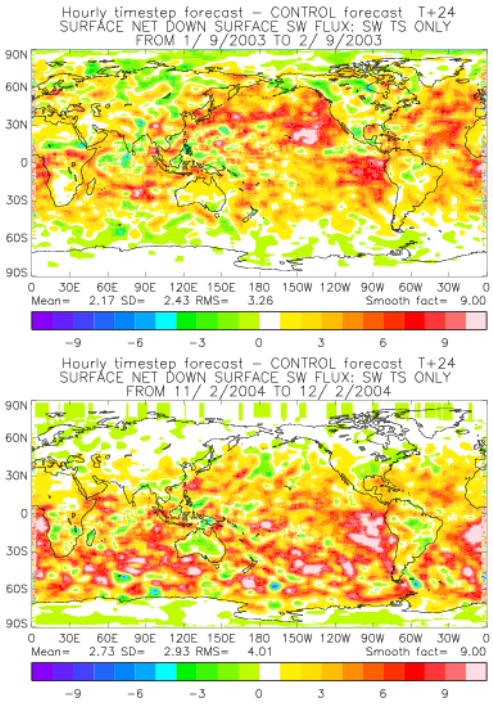

Figure 12: HOURLY_RAD−CONTROL: T+24 net surface SW from both JJAS (top panel) and DJF (bottom panel).

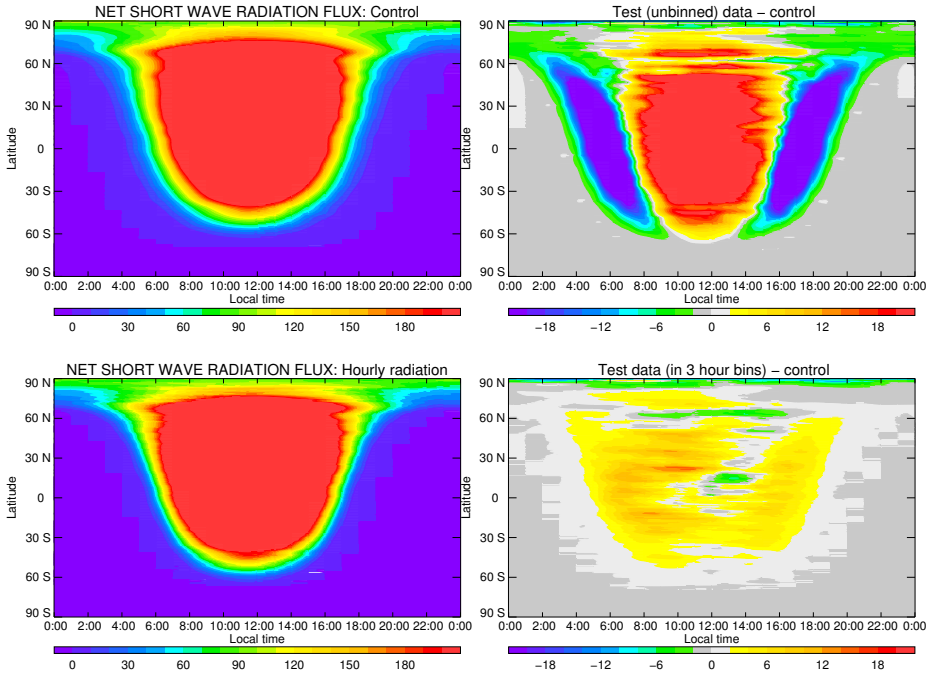

Figure 13: Zonally (local time) averaged SW surface flux.

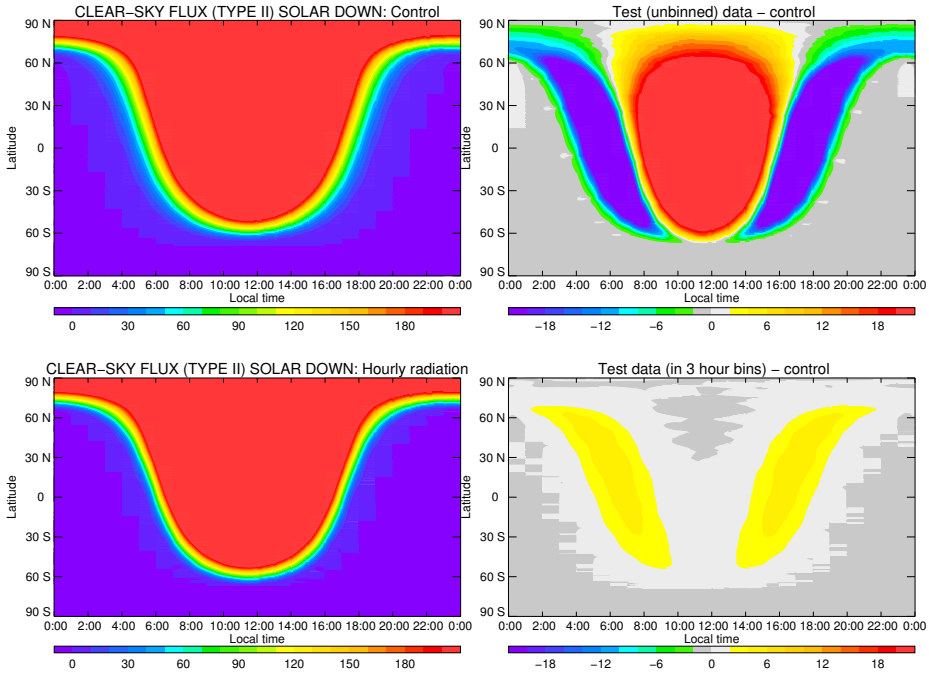

Figure 14: Zonally (local time) averaged clear-sky SW surface flux.

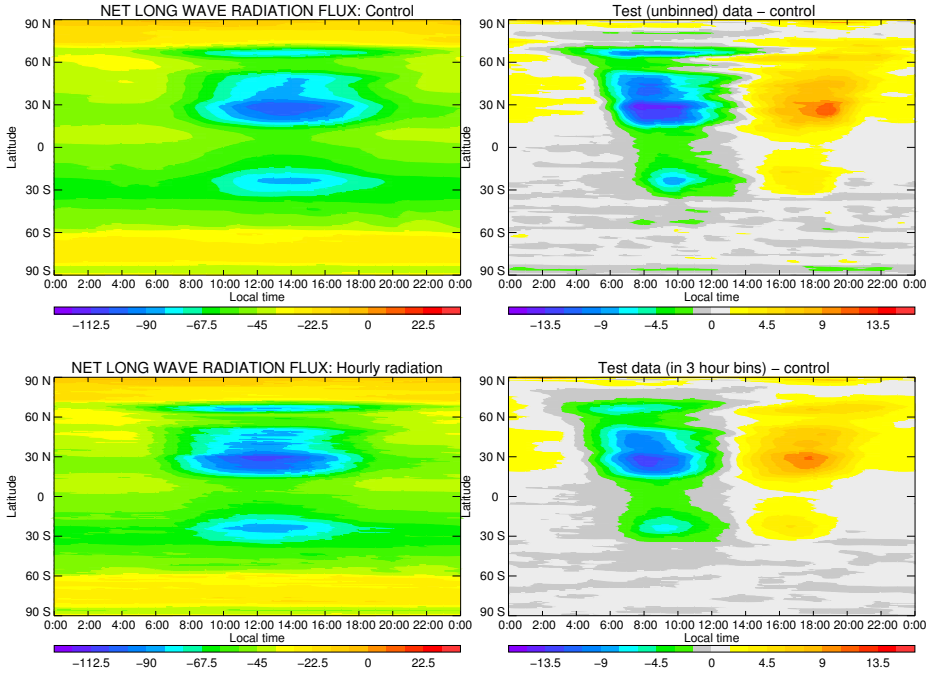

Figure 15: Zonally (local time) averaged LW surface flux.

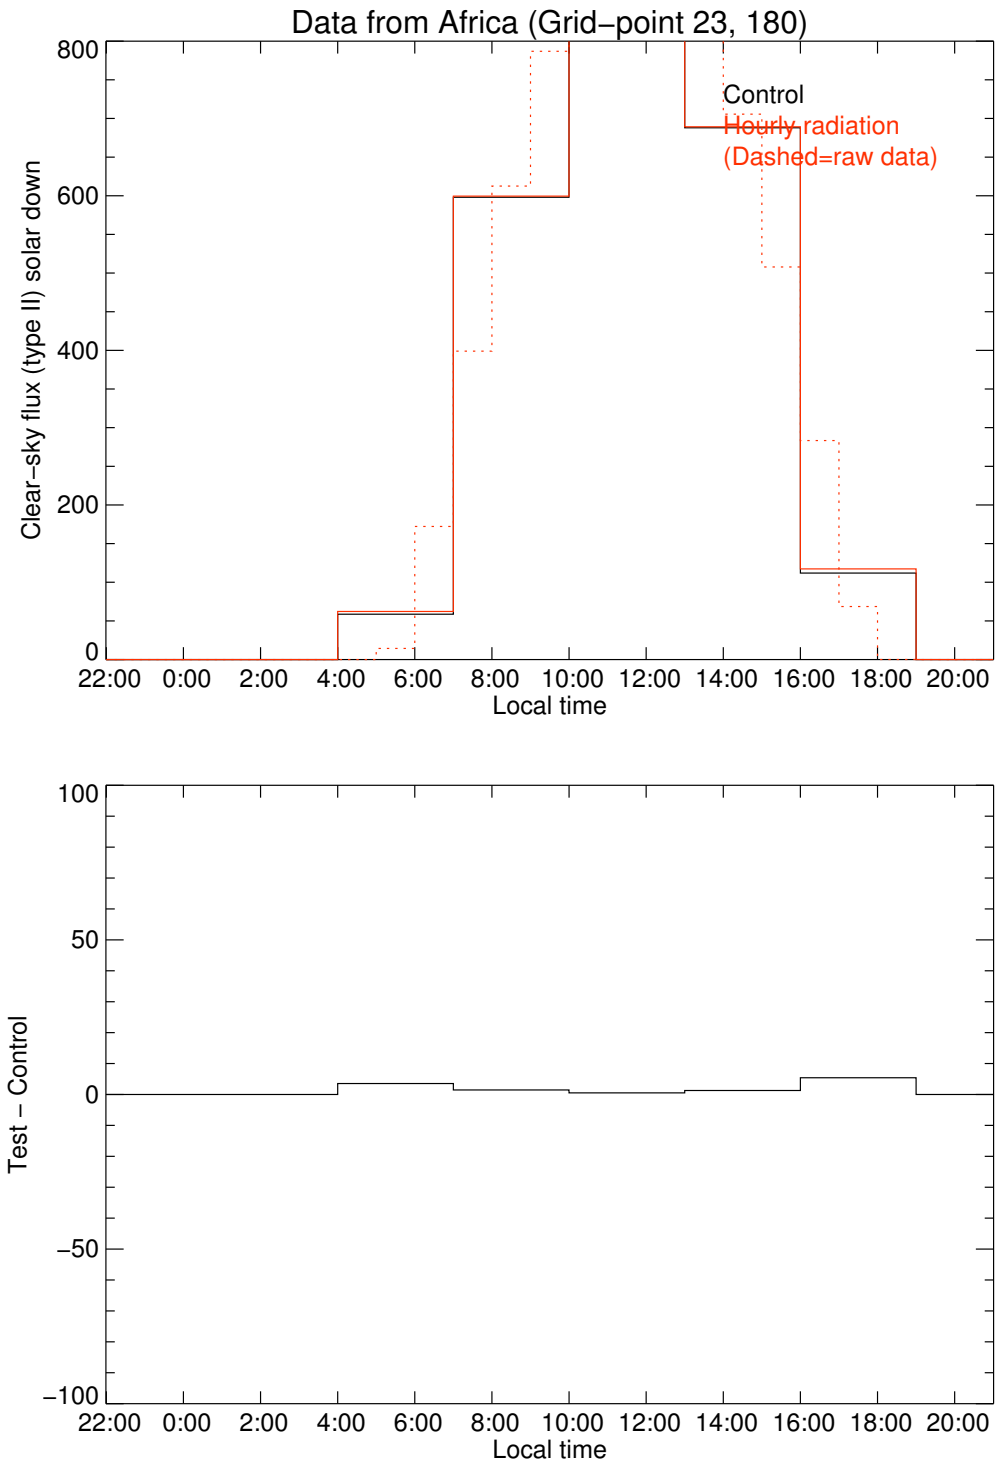

Figure 16: Clear-sky surface SW (down) over the Africa grid-point.

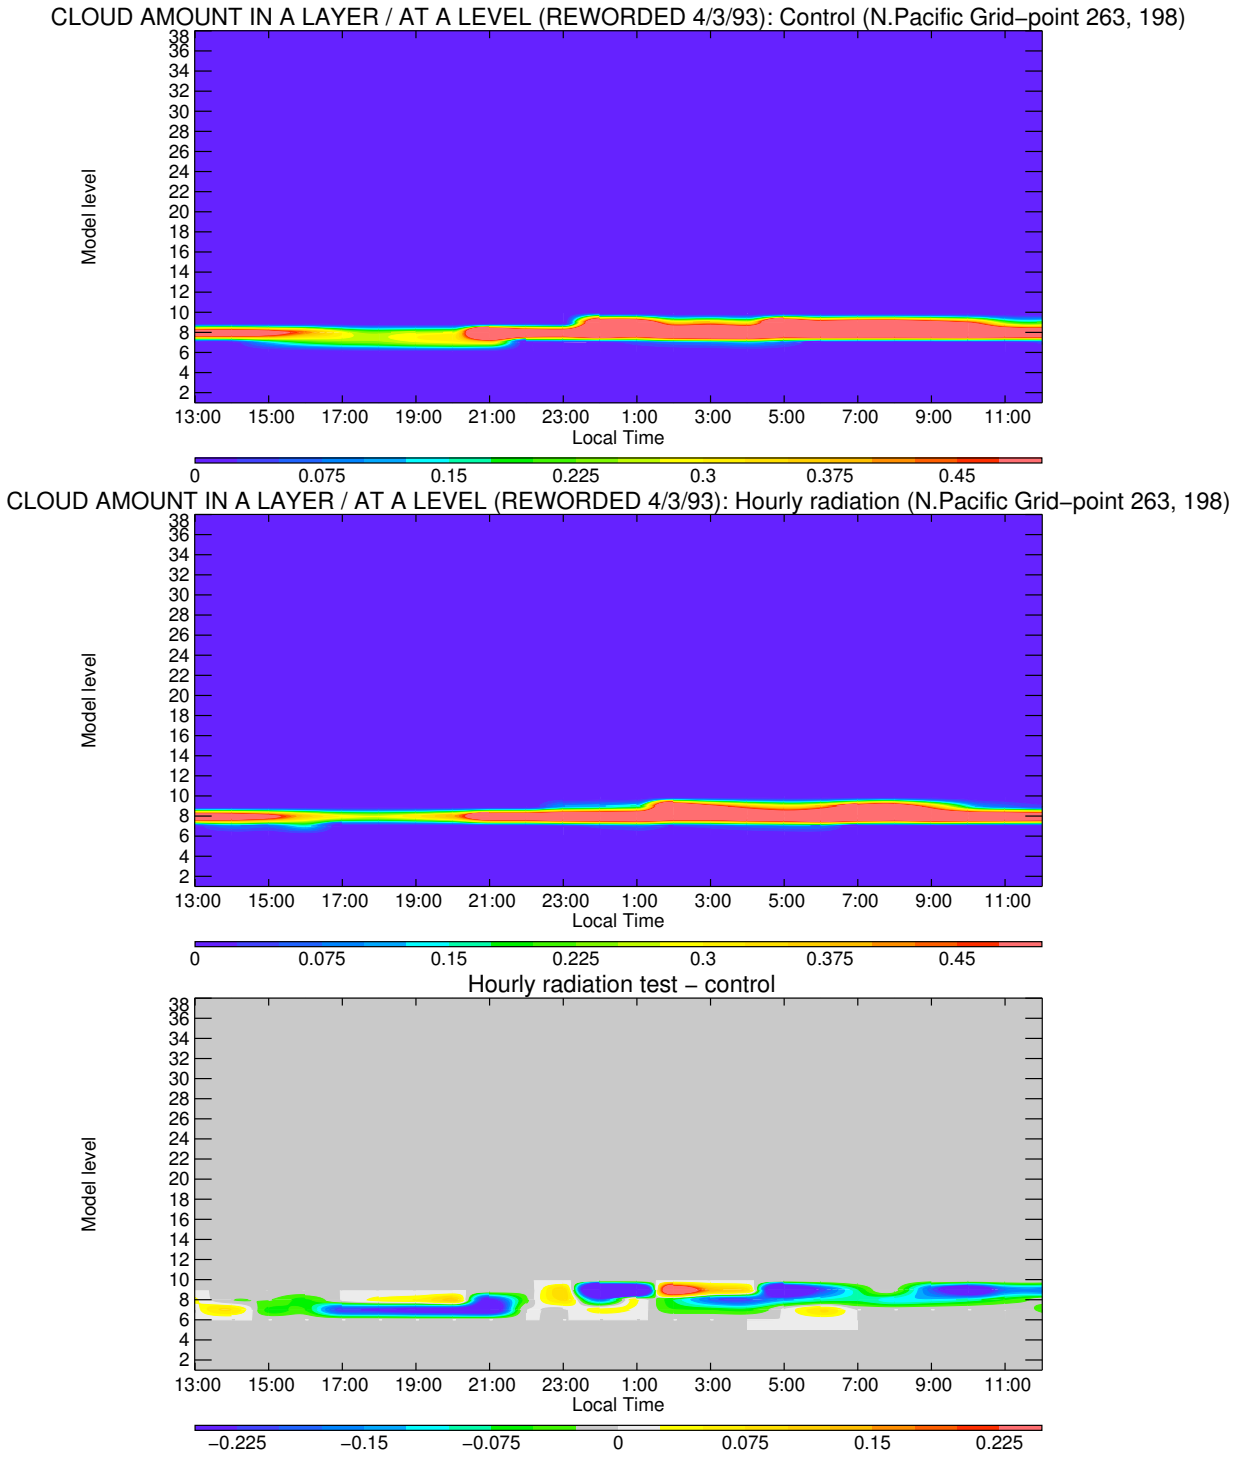

Figure 17: Bulk cloud amount over the N. Pacific grid-point.

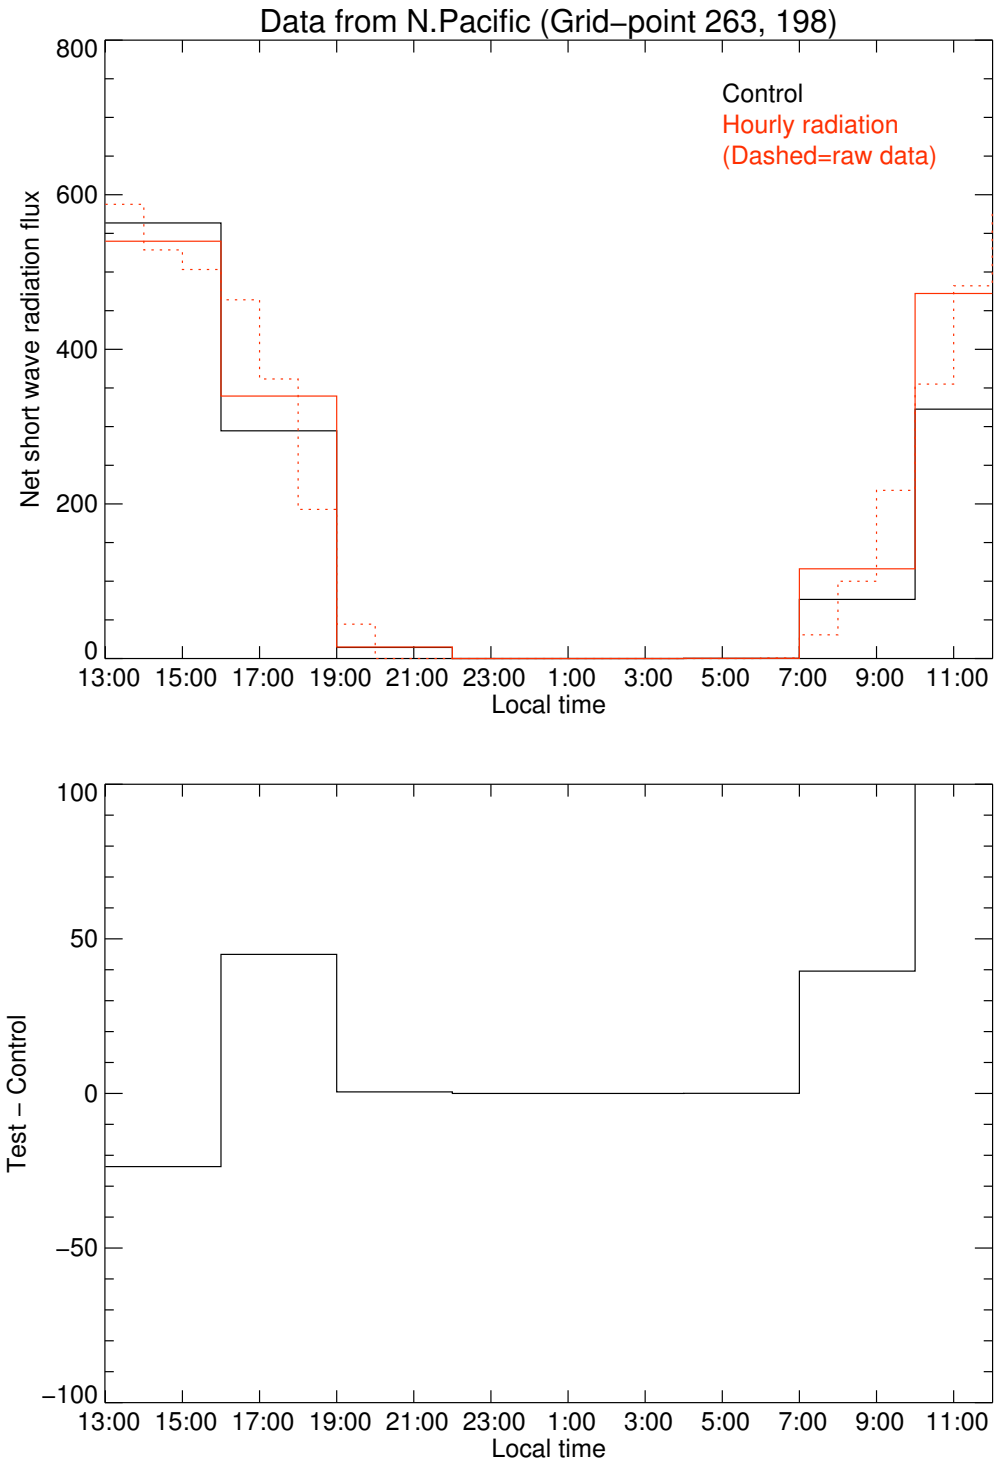

Figure 18: Net surface SW flux over the N. Pacific grid-point.

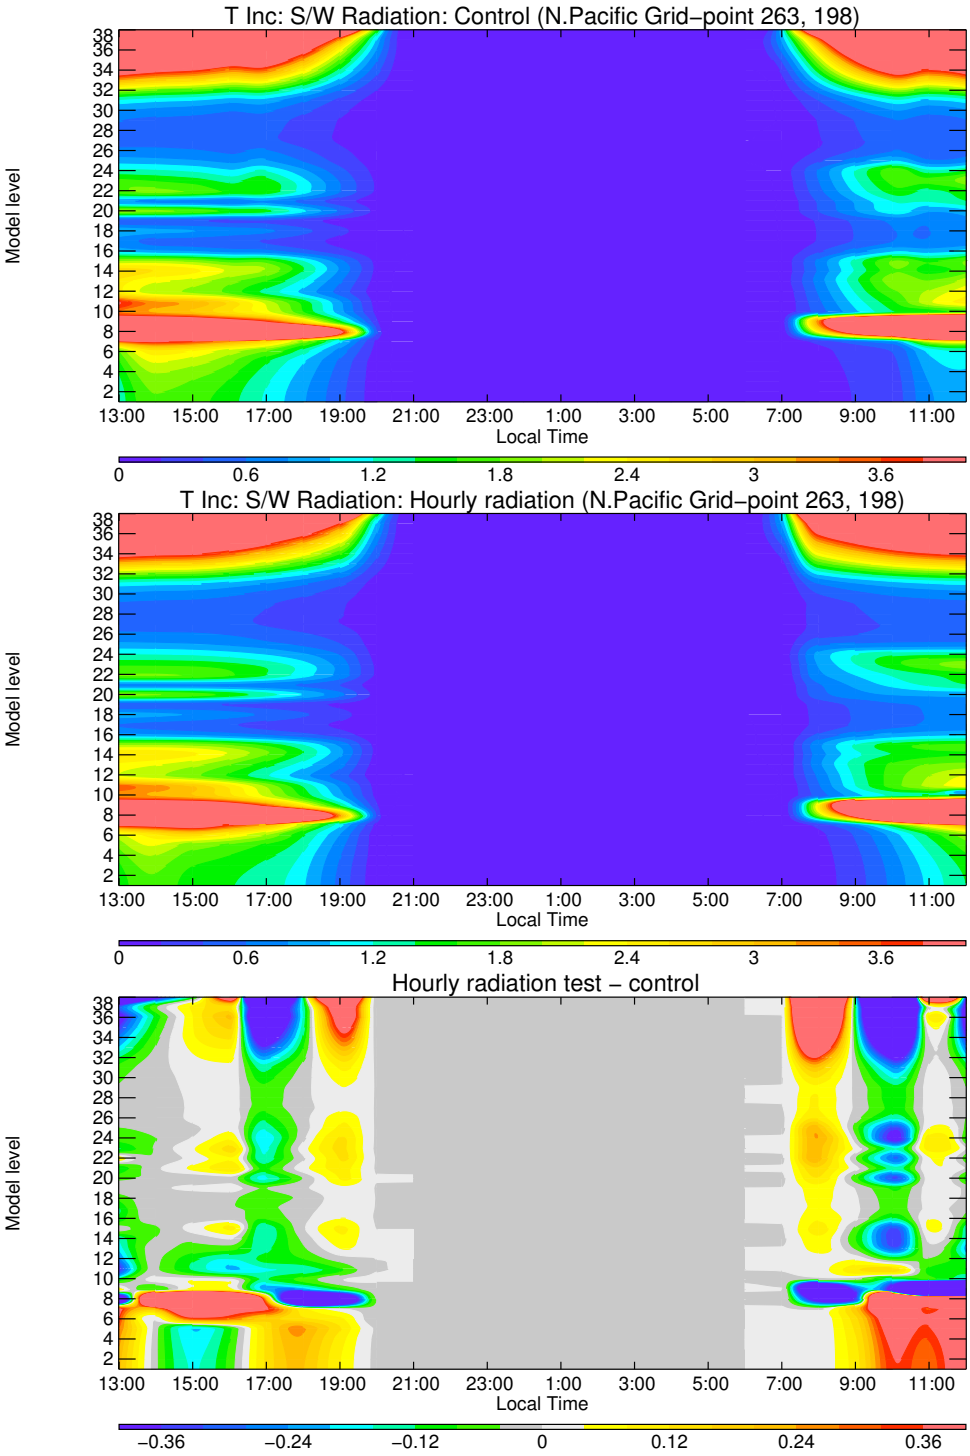

Figure 19: SW temperature increments (K/day) over the N. Pacific grid-point.